

Surfing the Digital Wave

Generating Personalised TV Listings using Collaborative, Case-Based Recommendation

Barry Smyth & Paul Cotter

Department of Computer Science
University College Dublin
Belfield, Dublin 4, Ireland
{Barry.Smyth, Paul.Cotter}@ucd.ie

Abstract. In the future digital TV will offer an unprecedented level of programme choice. We are told that this will lead to dramatic increases in viewer satisfaction as all viewing tastes are catered for all of the time. However, the reality may be somewhat different. We have not yet developed the tools to deal with this increased level of choice (for example, conventional TV guides will be virtually useless), and viewers will face a significant and frustrating information overload problem. This paper describes a solution in the form of the PTV system. PTV employs user profiling and information filtering techniques to generate web-based TV viewing guides that are personalised for the viewing preferences of individual users. The paper explains how PTV constructs graded user profiles to drive a hybrid recommendation technique, combining case-based and collaborative information filtering methods. The results of an extensive empirical study to evaluate the quality of PTV's case-based and collaborative filtering strategies are also described.

1 Introduction

With the advent of new cable and satellite services, and the next generation of digital TV systems, we will soon be faced with an unprecedented level of programme choice. Where we have tens of TV channels today, tomorrow we will have hundreds, and soon after that it will be thousands. If we believe the hype, we are entering a new age of television viewing, an age of incredible choice and unprecedented viewing satisfaction. However, while increased programme choice does offer the *potential* for improved viewing satisfaction, the reality may be somewhat different. We have not yet developed the tools to deal with this new level of choice, and it will become increasingly difficult to find out what programmes are on in a given week, never mind locating a small set of relevant programmes for a quiet evening's viewing.

Consider for example the traditional TV guide, listing programming information on local channels for up to a week in advance. The days of a slim, easy to digest 30 page volume are essentially gone. Looking to the US for a sign of things to come we

notice, with some consternation, that the current issue of the TV Guide (that weekly bible for American channel surfers) runs to nearly 400 pages of indigestible schedule charts. Moreover, the way that we interact with our TV sets will also have to change. Those rapid "remote-controlled surfs", that prove so effective (and so irritating to your partner) for 10 or 20 channels, will no longer be a viable means of finding out what is on at a given time. A 10 second per channel surf over even a modest 200 channel service will take about 35 minutes to complete! The digital TV vendors do recognise this as a serious information overload problem, and in response they are now offering electronic programme guides to help users to navigate this digital maze. However, these guides are relatively crude and offer little more than a static category based view of the evenings programming; the burden of search remains with the user.

This paper describes the PTV system (<http://ptv.ucd.ie>), which offers a working solution to the problem of locating relevant programme information quickly and easily. PTV combines user profiling and case-based reasoning (CBR) techniques to generate electronic TV viewing guides that are carefully personalised for the viewing preferences of individual users ([2, 6, 7]). At the present time, these electronic guides are Web based and delivered over the Internet to desktop PCs, but of course the advent of WebTV and cable-internet services will allow PTV to deliver personalised programme information directly to the TV set.

The remainder of this paper is organised in the following way. The next section provides an overview of the PTV system, describing its various sources of knowledge and main functional components. Section 3 focuses on PTV's user profiling and case-based recommendation strategy. Before concluding, section 4 reports the results of an extensive empirical study to evaluate the quality of the personalised programme guides that are produced, and the effectiveness of the case-based and collaborative recommendation strategies. Finally, a new appendix has been added to indicate the current state of the PTV system including the results of a recent online survey that add further support to the PTV concept and, we believe, paves the way for the use of collaborative, case-based recommendation methods in a wide range of personalised media service in the future.

2 The PTV System

PTV is a client-server system operating over the Web, allowing users to register, login, and view their personalised TV guides as specially customised Web pages. The architecture of PTV is shown in Figure 1. A standard Web browser provides the required client functionality, and all user interaction is handled via the HTML Forms interface. The heart of PTV lies with its server-side components, which handle all the main information processing functions such as user registration and authentication, user profiling, guide compilation, and the all-important programme recommendation and grading.

In the following sections we will concentrate on user profiling in PTV, focusing on how these profiles are used to deliver personalised content and, in particular, how PTV can make intelligent recommendations to PTV subscribers. However, in this

section we will provide a suitable backdrop for these future discussions by taking a broad look at the form and function of PTV's main components.

Profile Database & Profiler: The key to PTV's personalisation facility is an accurate database of user profiles. Each user profile encodes the TV preferences of a given user, listing channel information, preferred viewing time, liked and disliked programmes, subject preferences, etc (see Figure 1). Preliminary profile information is collected from the user at registration time in order to bootstrap the personalisation process. However, the majority of information is learned from grading feedback provided by the user; each recommended programme is accompanied with grading icons allowing the user to explicitly evaluate the proposed recommendation (see also section 3.1).

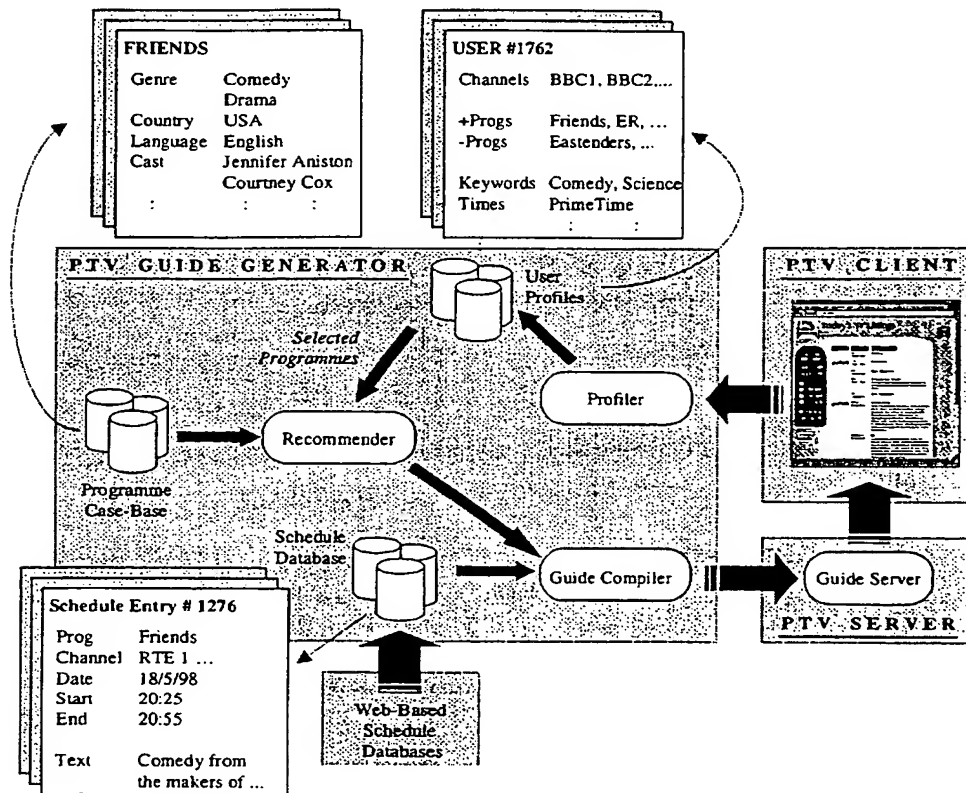


Fig. 1. An overview of the PTV system.

Programme Case-Base: This database contains the programme content descriptions (programme cases). Each entry describes a particular programme using features such as the programme title, genre information, the creator and director, cast or presenters, the country of origin, and the language; an example programme case for the comedy

'Friends' is shown in Figure 1. This information repository is crucial for the case-based recommendation component of PTV (see Section 3.2).

Schedule Database: This database contains TV listings for all supported channels. Each listing entry includes details such as the programme name, the viewing channel, the start and end time, and typically some text describing the programme in question (see the schedule entry example in Figure 1). The schedule database is constructed automatically from online schedule resources (e.g., online teletext pages and static entertainment guides) by PTV's schedule agents. Each agent is designed to mine a particular online resource for relevant schedule information and the results of these many parallel searches is the compilation of a rich schedule database.

Recommender: The recommendation component is the intelligent core of PTV. Its job is to take user profile information and to select new programmes for recommendation to a user. In the next section we will explain how PTV uses a hybrid recommendation approach that combines case-based and collaborative recommendation strategies (see section 3.2 and 3.3).

Guide Compiler: To compile a personalised TV guide for a given date and user, PTV constructs two programme lists: (1) a list consisting of those programmes listed as positive in the user's profile, along with those programmes selected for recommendation (which of course do not occur in the profile); (2) a list of all programmes to be aired on the specified date by a channel listed in the user's profile. The intersection of these two lists is the set of programmes that will be used to compile the user's personalised guide. The guide itself is a HTML page that is dynamically produced by drawing on programme and schedule information from the appropriate databases.

3 A Hybrid Information Filter

Ultimately the success of PTV will be measured in terms of the quality of its personalised guides, and this in turn depends on the quality of the user profiles and recommendation strategies that drive the guide compilation process. PTV harnesses two recommendation strategies to base its recommendations on the programmes that a given user has liked in the past (case-based or content-based) and on the programmes that similar users like (collaborative). In this section we look at PTV's profiling and recommendation components in more detail.

3.1 User Profiling for Programme Recommendation

In PTV each user profile contains two types of information, domain preferences and programme preferences. The former describe general user preferences such as a list of available channels, preferred viewing times, subject keywords, in addition to guide preferences such as whether guide programmes are to be sorted according to viewing time or channel. Programme preferences are represented as two lists of programme titles, a positive list containing programmes that the user has liked in the past, and a negative list containing programmes that the user has disliked.

At registration time a new user is invited to provide basic information including domain and programme preferences. This initial profile is needed to bootstrap the recommendation process, but usually only constitutes a restricted snapshot of a user's preferences. The left-hand screen shot of Figure 2 shows part of the user profile input screen used to gather explicit user information during registration time; indeed users can also use this facility to display and manually edit their own profile.

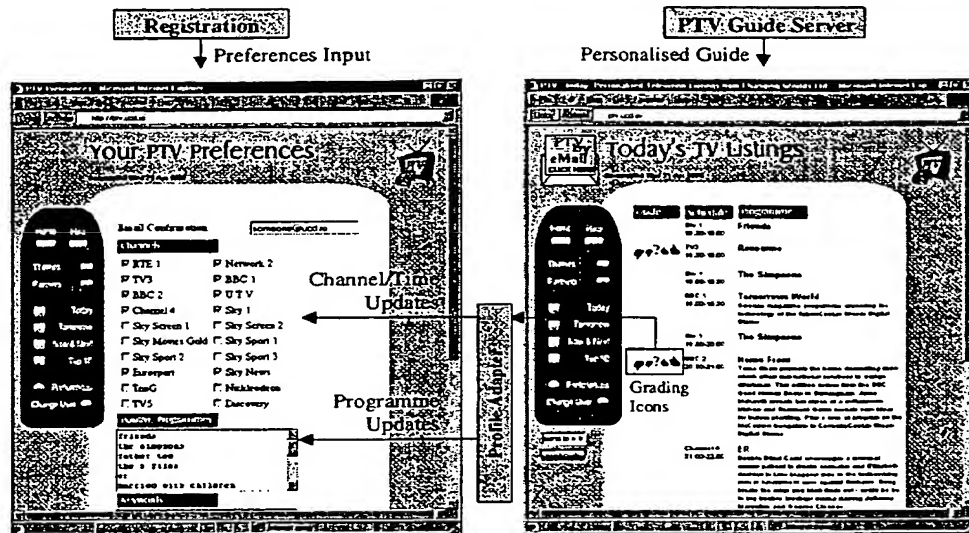


Fig. 2. User profiles and feedback

Of course while manual profile editing has its advantages (usually in terms of profile accuracy) it is a burden for users. In particular, we have found that users are happy to provide complete domain preferences but tend to provide only limited programme preferences. For this reason, PTV includes an automatic profile update facility that is driven by direct user feedback through a set of grading icons listed beside recommended guide programmes. PTV's profiler can use this feedback information to automatically alter a user's profile in a number of ways. The simplest modification is to update the programme preference lists by adding positively or negatively graded programmes to the appropriate list. However, the domain preferences can also be altered. For example, viewing time preferences can be adjusted if a user frequently prefers prime-time programmes to morning shows. This long-term feedback connection between user and system is vital if PTV is to maintain an accurate picture of each user over time.

3.2 Case-Based Recommendation

The basic philosophy in case-based recommendation is to recommend items that are similar to those items that the user has liked in the past (see also [1, 5, 11]). For PTV, this means recommending programmes that are similar to the programmes in the positive programme list and dissimilar to those in the negative programme list. Three components are needed for case-based-recommendation: (1) content descriptions for all TV programmes (see the programme case-base in section 2 and Figure 1); (2) a compatible content description of each user's profile; (3) a procedure for measuring the similarity between a programme and a user.

PTV's programme case-base has already been outlined in section 2 and an example case is shown in Figure 1. Each case is described as a set of features and the similarity between two cases can be defined as the weight sum of the similarity between corresponding case features. However, there is no direct means of computing the similarity between a case and a user profile, as user profiles are not described as a set of case features. Instead each raw user profile is converted into a feature-based representation called a *profile schema*. Basically, the profile schema corresponds to a content summary of the programme preferences contained in a user profile, encoded in the same features as the programme cases. The similarity between a profile and a given programme case can then be computed using the standard weighted-sum similarity metric as shown in equation 1; Where $f_i^{\text{Schema}(u)}$ and f_i^p are the i^{th} features of the schema and the programme case respectively.

$$1. \text{Pr gSim}(\text{Schema}(u), p) = \sum w_i \bullet \text{sim}(f_i^{\text{Schema}(u)}, f_i^p)$$

The main problem associated with case-based methods is the knowledge-engineering effort required to develop case representations and sophisticated similarity models. In addition, because case-based methods make recommendations based on item similarity, the newly recommended items tend to be similar to the past items leading to reduced diversity. In the TV domain this can be a significant problem as we may find that all a user's recommendations are, for example, comedies if the majority of profile programmes are comedies.

3.3 Collaborative Recommendation

Collaborative recommendation methods such as automated collaborated filtering are an alternative to case-based techniques. Instead of recommending new programmes that are similar to the ones that the user has liked in the past, we recommend programmes that other *similar users* have liked ([1, 3, 4, 8, 9, 10]). Rather than compute the similarity between items, we compute the similarity between users, or more precisely the similarity between user profiles. Note that we have opted for a lazy-approach to collaborative filtering rather than the more traditional eager approach where the user-base is pre-processed in to virtual communities prior to recommendation. So the recommendations for a target user are based on the viewing preferences of k similar users.

PTV computes user similarity by using a simple graded difference metric shown in equation 2; where $p(u)$ and $p(u')$ are the ranked programmes in each user's profile, and $r(p_i^u)$ is the rank of programme p_i in profile u . The possible grades range from -2 to +2 and missing programmes are given a default grade of 0. Of course this is just one possible similarity technique that has proved useful in PTV, and any number of techniques could have been used, for example statistical correlation techniques such as Pearson's correlation coefficient (see eg., [3, 10]).

$$2. \text{PrfSim}(u, u') = \frac{\sum_{p \in p(u) \cup p(u')} |r(p_i^u) - r(p_i^{u'})|}{4 \cdot |p(u) \cup p(u')|}$$

$$3. \text{PrgRank}(p, u) = \sum_{u' \in U} \text{PrfSim}(u, u')$$

Once PTV has selected k similar user profiles for a given target user, a recommendation list is formed from the programmes in these similar profiles that are absent from the target profile. This list is then ranked and the top r programmes are selected for recommendation. The ranking metric is shown in equation 3; U is the subset of k nearest profiles to the target that contain a programme p . This metric biases programmes according to their frequency in the similar profiles and the similarity of their recommending user. In this way popular programmes that are suggested by very similar users tend to be recommended.

Collaborative filtering is a powerful technique that solves many of the problems associated with case-based methods. For example, there is no need for content descriptions or sophisticated case similarity metrics. In fact, high quality recommendations, that would ordinarily demand a rich content representation, are possible. Moreover, recommendation diversity is maintained as relevant items that are dissimilar to the items in a user profile can be suggested.

Collaborative filtering does suffer from a number of shortcomings. There is a startup cost associated with gathering enough profile information to make accurate user similarity measurements. There is also a latency problem in that new items will not be recommended until these items have found their way into sufficiently many user profiles. This is particularly problematic in the TV domain because new and one-off programmes occur regularly and do need to be considered for recommendation even though these programmes will not have made it into any user profiles.

The key to PTV's success is the use of a combined recommendation approach. For a given guide, a selection of programmes is suggested, some are case-based recommendations (including new or one-off programmes) while others are collaborative recommendations. In particular, recommendation diversity is ensured through the use of collaborative filtering and the latency problem can be solved by using case-based methods to recommend new or one-off programmes.

4 Experimental Studies

PTV's normal mode of operation involves the generation of daily personalised TV guides containing a list of programme recommendations predicted to be of interest to each user. Guides contain, on average, 3 new recommendations per day. The hope is that all of these recommendations will be relevant, but of course the reality will inevitably be somewhat different. In this experiment we look at the gradings that users provided for the programme recommendations that they received in their daily guides. The primary question to be answered is whether or not users consider PTV's recommendations to be useful; that is, how often are the recommendations graded as relevant? In addition we are interested in comparing the recommendation quality of collaborative and case-based techniques.

4.1 Setup

The following results are based on an online evaluation by PTV users (mostly students and staff from University College Dublin and Trinity College Dublin) during March 1998. At this time the PTV system contained a population of approximately 200 users and a case-base of 400 programmes, which provided about 30% coverage of a typical week of television, and about 60% coverage of the prime-time viewing slots. During the experimental period a total of 2000 individual programme guides were requested. Each guide contained 3 new programme suggestions generated using either a collaborative approach or a case-based approach. This allows us to independently assess the relative competence of the collaborative filtering and case-based approaches. In addition, we generated guides by picking programme suggestions at random. These guides provide us with a basic benchmark against which to judge the quality of our two recommendation strategies.

4.2 Method

Each time a new programme is recommended as part of a personalised guide, the user is invited to grade the recommendation on a 5 point scale ranging from -2 (terrible) through 0 (no comment) to +2 (an excellent suggestion). These gradings are the raw data for the experiment; each grading encodes the username, the programme name, the date, the grade itself, and whether the recommendation was a case-based one or a collaborative filtering one. Over the experimental period a total of approximately 1000 individual gradings were saved from 100 different users; on average each user was submitting 10 grades over the test period. This provided us with three sets of data: (1) the grades associated with case-based recommendations; (2) the grades associated with collaborative filtering recommendations; (3) the grades associated with random recommendations.

4.3 Results

The quality of PTV's suggestions could be measured in any number of ways. For example, we could simply calculate the average grade assigned to collaborative, case-based, and random recommendations. However, we feel a better approach is to look at recommendation quality in the context of individually generated guides. For this reason our approach is to take each guide in turn, and to count the percentage of users that received 'n' or more good recommendations per day. This allows us to count the number of users that received at least 1, 2, or 3 good recommendations per day, which to our minds is a far better way of evaluating recommendation success in this setting. The results are displayed in Figure 3. Clearly, the results for PTV are very positive with 96% and 78% of users receiving at least one good new programme suggestion per day, depending on whether the guide was generated using collaborative filtering techniques or case-based methods. By comparison only 27% of users found one of the random recommendations to be worth watching. In fact, the results show that PTV recommended 3 good programmes a day more often than the random method recommended 1 good programme (42% or 32% versus 27%).




Fig. 3. PTV Guide Quality Results

TV programme recommendation is a difficult task. Viewers display extremely varied, and often inconsistent, viewing preferences and patterns. In addition, recommendations for a given guide on a given day are drawn from a limited set. There may be very few programmes scheduled for a particular day that suit a given user. Therefore, the results presented here bode extremely well for PTV's long-term recommendation prospects. Both collaborative and case-based methods have been shown to perform extremely well, and form an ideal partnership in PTV to ensure the recommendation of a diverse range of regular and new programmes.

These experiments represent an initial attempt at evaluation and admittedly the use of random recommendations as an evaluation benchmark is quite limited – it is easy to imagine more sophisticated benchmarks. However, in the Appendix we provide additional support in favour of PTV compiled from an extended user trial during March 1999.

5 Conclusions

As the latest satellite and digital TV services beam hundreds (and soon, thousands) of TV channels into our homes, we are faced with a significant choice problem, and the job of finding the right TV programme at the right time becomes increasingly difficult. In fact, instead of witnessing an increase in viewer satisfaction, some commentators have predicted quite the opposite, as viewers fail to come to grips with the new range of channels and fall into channel-hopping oblivion. In this paper we have described one possible solution to the problem, a solution that involves automatically generating personalised TV guides for individuals based on their learned viewing preferences, each guide containing information about a select set of programmes that are relevant to a particular user.

We have explained how the PTV system uses a hybrid approach to recommendation, which combines collaborative and case-based techniques to make high quality and diverse programme recommendations, and encompassing the recommendation of new or one-off programmes as well as regular programmes. We have also presented an evaluation of the system to demonstrate the effectiveness of PTV's recommendation components and to examine the separate contributions of the collaborative filtering and the case-based strategies.


Recently the AI community has been challenged to solve the task of automatically producing dynamic and personalised Web content (IJCAI-1997 Challenge). PTV's hybrid recommendation approach represents a direct response to this challenge. The initial results in the TV domain suggest a promising future for this approach in a wide range of information filtering and personalisation tasks.

Appendix

Initially the PTV project started life as a demonstration of what we saw as an important idea, namely that, in general, Internet content could be effectively personalised for the needs of individuals, and more specifically that users could receive accurate personalised TV guides to help them cope with the onslaught of digital TV. Since going live, and without significant advertising or marketing, the PTV system has attracted and profiled over 5000 registered users (growing by 500 - 1000 users per month) and over 15,000 personalised TV guides per month.

The following charts summarise the relevant results of a recent user survey carried out during March 1999. In particular they indicate that the recommendation approach is producing good quality guides at an acceptable speed thus supporting the claim that

collaborative, case-based recommendation provides an effective and efficient personalisation technique.



References

1. Balabanovic M., Shoham Y.: FAB: Content-Based Collaborative Recommender. Communications of the ACM, 40(3) (1997) 66-72
2. Baudisch, P.: Recommending TV Programs: How far can we get at zero effort. In: Proceedings of the AAAI-98 Workshop on Recommender Systems, Wisconsin, USA, (1998) 16-19
3. Billsus, D. & Pazzani, M. J.: Learning collaborative Information Filters. In: Proceedings of the International Conference on Machine Learning, Wisconsin, USA, (1998)
4. Goldberg D., Nichols D., Oki B. M., Terry D.: Using Collaborative Filtering to Weave an Information Tapestry. Communications of the ACM, 35(12) (1992) 61-70
5. Hammond, K. J., Burke, R., and Schmitt, K.: A Case-Based Approach to Knowledge Navigation. In: (Leake, D.B, ed.) Case-Based Reasoning Experiences Lessons and Future Directions, MIT Press, (1996) 125-136
6. Jennings, A. & Higuchi, H.: A user model neural network for a personal news service. User Modeling and User-Adapted Information, 3(1) (1993).1-25
7. Kay J.: Vive la Difference! Individualised Interaction with Users. In: Proceedings IJCAI '95, Montréal, Canada, (1995). 978-984
8. Konstan J. A., Miller B. N., Maltz D., Herlocker J. L., Gordan L. R., Riedl J.: GroupLens: Applying Collaborative Filtering to Usenet News. Communications of the ACM. 40(3) (1997) 77-87
9. Maltz D., Ehrlich K.: Pointing the Way: Active Collaborative Filtering. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '95) ACM Press, New York, N.Y., (1995). 202-209
10. Shardanand, U. & Maes, P.: Social Information Filtering: Algorithms for Automating 'Word of Mouth'. In: Proceedings of the Conference on Human Factors in Computing Systems (CHI95), ACM Press, New York, N.Y., (1995). 210-217
11. Watson, I., Applying Case-Based Reasoning: Techniques for Enterprise Systems, Morgan-Kaufmann, (1997)

THIS PAGE BLANK (USPTO)